Chapter 1

MULTI-AGENT CONTRACT NEGOTIATION

Peyman Faratin
MIT Sloan School of Management
peyman@mit.edu

Abstract Two computational decision models are presented for the problem of de-centralized contracting of multi-dimensional services and goods between autonomous agents. The assumption of the models is that agents are bounded in both information and computation. Heuristic and approximate solution techniques from Artificial Intelligence are used for the design of decision mechanism that approach mutual selection of efficient contracts.

Keywords: Negotiation, Multi-Agent Systems, Game Theory, Artificial Intelligence.

Introduction

The problem of interest in this chapter is how autonomous computational agents can approach an efficient trading of multi-dimensional services or goods under assumptions of bounded rationality. Trading is assumed to involve negotiation, a resolution mechanism for conflicting *preferences* between selfish agents. We restrict ourselves to a monopolistic economy of two trading agents that meet only once to exchange goods and services. Agents are assumed to be bounded in both information and computation. Information needed for decision making is assumed to be bounded due to both external and internal factors, social and local information respectively. Agents have limited social information because they are assumed to be selfish, sharing little or no information. In addition to this agents may also have limited *local* information (for example over their own preferences) because of complexity of their local task(s). Computation, in turn, is a problem in contract negotiation because of the combinatorics of *scale*. Computation is informally defined as the process of searching a space of possibilities [11]. For a contract with 100 issues and only two alternatives for each issue, the size of the search space is roughly 10^{30} possible contracts, too large to be explored exhaustively.

The unbounded formulation of such an economical problem has long been the central concern of classic game theory which has produced a number of models of social choice. For this reason game theory models have become strong candidates for models of social agents. Surprisingly, such apparently simple games can be used to conceptualize a variety of synthetic, meaningful and formal prototypical context as games. Therefore, such models can be used to design and engineer multi-agent systems as well as analyze the behaviour of the resulting social artifact using the logical tools of the models. However, the underlying unbounded assumptions of classic game theory is problematic for the design of computational systems [2].

Artificial Intelligence (AI) on the other hand has long considered models of the relationship between knowledge, computation and the quality of solution (henceforth referred to as the K-C-Q relationship) [7]. AI has shown that there exists a hierarchy of tradeoffs between K, C and Q, with models that achieve perfect optimal results (like game theory models) but at the cost of requiring omniscience and unbounded agents, to models that sacrifice optimality of Q for a more realistic set of requirements over K and C [12]. Different agent architectures are then entailed from different K-C-Q relationship theories.

In the next two sections two such computational models of negotiation are proposed, one deductive and the other agent-based simulation, that can be analyzed as two different games. The aim of these models has been to attempt to address some of the computational and knowledge problems mentioned above. In particular, in the first model the types of problems of interest is when K is limited because agents have at best imperfect and at worst no knowledge of the others' utility functions. The best an agent can do is to reason with imperfect knowledge by forming approximations of others' utilities. In the second model the knowledge problem is even more extensive because agents in addition are assumed to have an incomplete knowledge of their *own* utility functions.

A Bargaining Game

In this model there are two players (a and b) representing one consumer and one producer of a service or a good. The goal of the two agents is to negotiate an outcome $x \in X$, where X is the set of possible contracts describing multi-dimensional goods/services such as the price of the service, the time at which it is required, the quality of the delivered service and the penalty to be paid for reneging on the agreement. If they reach an agreement, then they each receive a payoff dictated by their utility function, defined as $U_i : X \rightarrow [0, 1], i \in \{a, b\}$. If the agents fail to reach any deal, they each receive a conflict payoff c. However, from the set X, only a subset of outcomes are "reachable". Call the set of *feasible outcomes* B, containing those agreements that are *individually rational* and bounded by the *pareto optimal* line [13]. An agreement is individually rational if it assigns each agent a utility that is at least as large as the agent can guarantee for itself from the conflict outcome $\mathbf{x_c}$. Pareto optimality is informally defined as the set of outcomes that are better for *both* agents [1]. It is often used as a measure of the efficiency of the social outcome. Given the game $(B, \mathbf{x_c})$, the protocol, or *"rules of encounter"* [8], normatively specifies the *process* of negotiation. The protocol chosen for this game is the alternating sequential model in which the agents take turns to make offers and counter offers [10]. The protocol terminates when the agents come to an agreement or time limits are reached or, alternatively, when one of the agents withdraws from the negotiation. This distributed, iterative and finite protocol was selected because it is un-mediated, supports belief update and places time bounds on the computational resources that can be utilized.

However, like chess for example, agents can have different negotiation strategies given the normative rules of the game. Two heuristic distributed and autonomous search strategies have been developed whose design has been motivated by the knowledge and computation boundedness arguments given above. One parametric mechanism, the *responsive mechanism*, is a mechanism that conditions the decisions of the agent directly to its environment such as the concessionary behaviour of the other party, the time elapsed in negotiation, the resources used, etc. [3]. However, the mechanism is known to have several limitations [4]. In some cases agents fail to make agreements, even though there are potential solutions, because they fail to explore different possible value combinations for the negotiation issues. For instance, a contract may exist in which the service consumer offers to pay a higher price for a service if it is delivered sooner. This contract may be of equal value to the consumer as one that has a lower price and is delivered later. However from the service provider's point of view, the former may be acceptable and the latter may not. The responsive mechanism does not allow the agents to explore for such possibilities because it treats each issue independently and only allows agents to concede on issues.

A second mechanism, called the *trade-off mechanism*, was developed to address the above limitations and consequently select solutions that lie closer to the pareto-optimal line, again in the presence of limited knowledge and computational boundedness [4]. Intuitively, a trade-off is where one party lowers its utility on some negotiation issues and simultaneously demands more on others while maintaining a constant overall contract utility. This, in turn, should make agreement more likely and increase the efficiency of the contracts. An algorithm has been developed that enables agents to make trade-offs between both quantitative and qualitative negotiation issues, in the presence of information uncertainty and resource boundedness for multi-dimensional goods [4]. The algorithm computes n dimensional trade-offs using techniques from fuzzy sim-

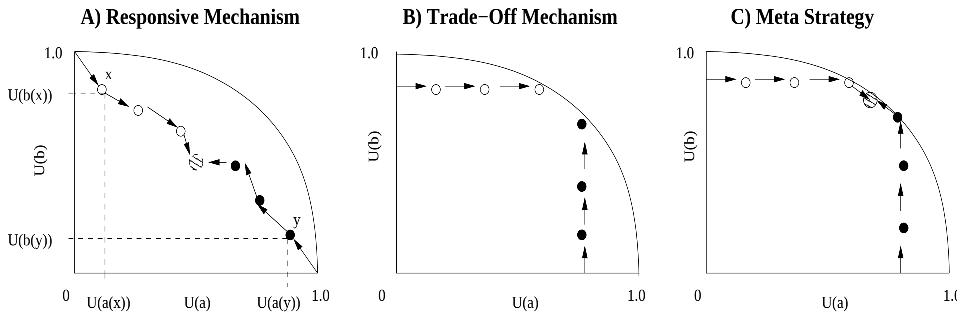

Figure 1.1. Utility Dynamics of the Mechanisms

ilarity [14] to approximate the preference structure of the negotiation opponent. It then uses a hill-climbing technique to explore the space of possible contract trade-offs for a contract that is most likely to be acceptable. The complexity of this algorithm has been shown to grow linearly with growing numbers of issues [4].

The details of the algorithms can be found in [3] and [4]. The dynamics of the contract utility generated by each of the above mechanisms and one possible combination is given in figure 1.1 A, B and C respectively for the alternating sequential protocol. The filled ovals are the utility of the offered contracts from agent a to agent b from agent a's perspective, and the unfilled ovals represent the utility of the offered contracts from agent b to agent a from agent b's perspective. The patterned oval represents the joint utility of the final outcomes. The pareto-optimal line is given by the curvilinear line connecting the two pairs of payoffs $(1, 0)$ and $(0, 1)$. Figure 1.1 A represents *a* possible execution trace where both agents generate contracts with the responsive mechanism. Each offer has lower utility for the agent who makes the offer, but relatively more utility for the other. This process continues until one of the agents is satisfied ($U^a(x^t_{b\to a}) \geq U^a(x^{t'}_{a\to b})$), where $x^t_{b\to a}$ is the contract offered by agent b to a at time t. This termination criteria is referred to as the cross-over in utilities. The responsive mechanism can select different outcomes based on the rate of concession adopted for each issue (the angle of approach to the outcome point in figure 1.1 A).

Figure 1.1 B represents another possible utility execution trace where both agents now generate contracts with the trade-off mechanism. Now each offer has the same utility for the agent who makes the offer, but relatively more utility for the other (movement towards the *pareto-optimal* line). The trade-off mechanism searches for outcomes that are of the same utility to the agent, but which *may* result in a higher utility for the opponent. Once again, this is a simplification for purposes of the exposition—an offer generated by agent a may indeed have decreasing utility to agent b (arrow moving *away* from the *pareto-optimal* line) if the similarity function being used does not correctly induce the preferences of the other agent.

Finally, agents can combine the two mechanisms through a meta strategy (figure 1.1 C). One rationale for the use of a meta-strategy is reasoning about the costs and benefits of different search mechanisms. Another rationale, observable from the example shown in figure 1.1 B, is that because the local utility information is private agents can not make an interpersonal comparison of individual utilities in order to compute whether a pareto optimal solution has indeed been reached. In the absence of a mediator the lack of such global information means negotiation will fail to find a joint solution that is acceptable to both parties. In fact agents enter a loop of exchanging the same contract with one another. Figure 1.1 C shows a solution where both agents implement a responsive mechanism and concede utility. This concession may, as shown in figure 1.1 C, indeed satisfy the termination conditions of the trade-off mechanism where offers cross-over in utilities. Alternatively, agents may resume implementing a trade-off algorithm until such a cross-over is eventually reached or time limits are reached. In general, the evaluation of which search should be implemented is delegated to a meta-level reasoner whose decisions can be based on bounding factors such as the opponent's perceived strategy, the on-line cost of communication, the off-line cost of the search algorithm, the structure of the problem or the optimality of the search mechanism in terms of completeness (finding an agreement when one exists), the time and space complexity of the search mechanism, and the expected solution optimality of the mechanism when more than one agreement is feasible.

A Mediated Game

In the above model the issues being negotiated over are assumed to be independent, where the utility to an agent of a given issue choice is independent of what selections are made for other issues. The utility function that aggregates the individual utilities under this assumption is then taken to be linear. This assumption significantly simplifies the agents' local decision problem of what issue values to propose in order to optimize their local utility. Optimization of such a linear function is achieved by hillclimbing the utility gradient. However, real world contracts, are highly inter-dependent. When issue interdependencies exist, the utility function for the agents exhibits multiple local optima. Multi-optimality results in firstly a more extensive bounded rationality problem since not only is computation limited but now also both *local* and global knowledge are limited. Local knowledge is limited because the agent now has to know and optimize a much more complicated utility function. Secondly, a methodological change from deductive models to simulation studies is needed due to the complex non-linearities involved in the system. The solution to these problems are briefly outlined below in a model of negotiation that departs from the more deductive model outlined above [5].

In this model a contract x is an N dimensional boolean vector where $x_i \in \{-1, +1\}$, represents the presence or absence of a "contract clause" i. The contract search policy is encoded in the negotiation protocol. Because generating contract proposals locally is both knowledge and computationally expensive we adopt an indirect single text protocol between two agents by delegating the contract generation process to a centralized mediator [9]. A mediator proposes a contract x^t at time t. Each agent then votes to accept or reject x^t. If both vote to accept, the mediator iteratively mutates the contract x^t and generates x^{t+1}. If one or both agents vote to reject, a mutation of the most recent mutually acceptable contract is proposed instead. The process is continued until the utility values for both agents become stable (i.e. until none of the newly contract proposals offer any improvement in utility values for either agent). Note that this approach can straightforwardly be extended to N party (i.e. multi-lateral) negotiation. The utility of the contract to an agent is defined as the linear combination of all the pairwise influences between issues.

Two computationally inexpensive decision algorithms were evaluated in this protocol: a hillclimber and a simulated annealer. A hillclimber only accepts a contract if and only if the utility of the contract x increases monotonically when *an* issue is changed. However, this steepest ascend algorithm is known to be incapable of escaping local maxima of the utility function. The other decision algorithm is based on the knowledge that search success can be improved by adding thermal noise to this decision rule [6]. The policy of decreasing T with time is called simulated annealing [6]. Simulated annealing rule is known to reach utility equilibrium states when each issue is changed with a finite probability and time delays are negligible.

To evaluate these algorithms simulations were run again with two agents a and b. The contract length N was set to 100 (corresponding to a space of 2^{100}, or roughly 10^{30} possible contracts) where each bit was initialized to a value $\{-1, +1\}$ randomly with a uniform distribution. The initial temperature was set to 10 and decreased in steps of 0.1 to 0. Final average utilities were collected for 100 runs for each temperature decrement. The left figure in figure 1.2 shows the observed individual payoffs for tests examining the relationship of C-Q with local utility metric of Q. One observation is that if the other agent is a local hill-climber, an agent is then individually better off being a local hill-climber, but fares very badly as local annealer. If the other agent is an annealer, the agent fares well as an annealer but does even better as a hillclimber. The highest *social* welfare, however, is achieved when both agents are annealers.

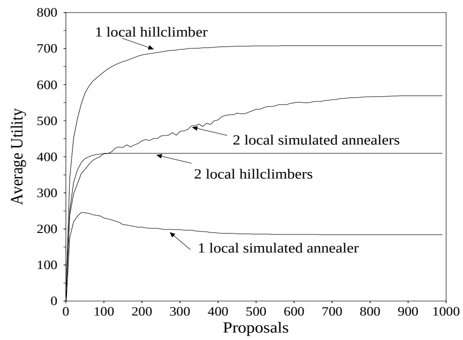

	Annealer	Hillclimber
Annealer	550/550	180/700
Hillclimber	700/180	400/400

Figure 1.2. Game Dynamics (left) and Final Payoff Matrix of the Game (right)

This pattern can be readily understood as follows. At high virtual temperature an annealer accepts almost all proposed contracts independently of the cost-benefit margins. Therefore, at high virtual temperature the simulated annealer is more explorative and "far sighted" because it assumes costs now are offset by gains later. This is in contrast to the myopic nature of the hillclimber where exploration is constrained by the monotonicity requirement. In the asymmetric interaction the cooperation of annealers *permits* more exploration of the contract space, and hence arrival to higher optima, of hillclimber's utility landscape. However, this cooperation is not reciprocated by hillclimbers who act selfishly. Therefore, gains of hillclimbers are achieved at the cost of the annealer. The right figure in figure 1.2 represents the underlying game as a matrix of final observed utilities for all the pairings of hillclimber and annealer strategies. The results confirm that this game is an instance of the prisoner's dilemma game [1], where for each agent the dominant strategy is hillclimbing. Therefore, the unique dominating strategy is for both agents to hillclimb. However, this unique dominating strategy is pareto-optimally dominated when both are annealers. In other words, the single Nash equilibria of this game (two hillclimbers) is the only solution not in the Pareto set.

Conclusions

The contracting problem was used to motivate two different heuristic and approximate agent decision models, both based on a realistic set of requirements over both K and C. However, the cost of these requirements is the sub-optimality of Q. This trade-off was demonstrated in both models by negotiation strategies selecting outcomes that are not pareto efficient. However, imperfections is a common feature of the world and real social systems have established personal and institutional mechanisms for dealing with such imperfections. Similarly, in future computational models are sought that are incremental, repeated and support feedback and error-correction. Learning and evolutionary techniques are two candidates for optimizing this trade-off given the environment of the agent.

References

[1] K. Binmore (1992). *Fun and Games: A Text on Game Theory* Lexington, Massachusetts: D.C. Heath and Company.

[2] P. Faratin (2000).*Automated Service Negotiation between Autonomous Computational Agents* Ph.D. Thesis, Department of Electronic Engineering, Queen Mary and Westfield College, University of London.

[3] Faratin, P and Sierra, C and Jennings, N.R (1998). *Negotiation Decision Functions for Autonomous Agents,* Journal of Robotics and Autonomous Systems, 24(3–4):159–182.

[4] Faratin, P and Sierra, C and Jennings, N.R (1998). *Using Similarity Criteria to Make Negotiation Trade-Offs,* International Conference on Multi-agent Systems (ICMAS-2000), Boston, MA., 119-126.

[5] P. Faratin, M. Klein, H. Samaya and Yaneer Bar-Yam (2001). *Simple Negotiating Agents in Complex Games: Emergent Equilibria and Dominance of Strategies,* In Proceedings of the 8th Int Workshop on Agent Theories, Architectures and Languages (ATAL-01), Seattle, USA.

[6] S. Kirkpatrick and C.D. Gelatt and M.P. Vecci (1983). *Optimization by Simulated Annealing*, Science 671–680.

[7] J. Pearl (1984). *Heuristics* Reading, MA: Addison-Wesley.

[8] J. S. Rosenschein and G. Zlotkin (1994). *Rules of Encounter* Cambridge, Massachusetts: MIT Press.

[9] H. Raiffa (1982). *The Art and Science of Negotiation*, Cambridge, MA: Harvard University Press.

[10] A. Rubinstein (1982). *Perfect equilibrium in a bargaining model*, Econometrica, 50:97–109.

[11] S. Russell and P. Norvig (1995). *Artificial Intelligence: A modern approach* Upper Saddle River, New Jersey: Prentice Hall.

[12] S. Russell and E. Wefald (1991). *Do the Right Thing* Cambridge, Massachusetts: MIT Press.

[13] J. von Neumann and O. Morgernstern (1944). *The Theory of Games and Economic Behaviour* Princeton, N.J: Princeton University Press.

[14] L. A. Zadeh (1971). *Similarity relations and fuzzy orderings* Information Sciences, 3:177–200.